

## **Level 2 Review Report**

September 2, 2002

A review of the Level 2 Decision system was held on August 1, 2002. The committee thanks the speakers for their hard work in preparing excellent talks.

The report is organized as follows: first are general comments on the Level 2 Decision system, followed by a summary of issues relevant to the Run 2A operation and the committee's recommendations for the near-term operations. Finally we discuss the proposals and recommendations for upgrading the Level 2 Decision system. The charge to the committee, membership of the committee, and review agenda are given in appendices. Slides from the talks can be found at [http://cdfsga.fnal.gov/upgrades/daq\\_trig/trigger/reviews/l2\\_review\\_020801.html](http://cdfsga.fnal.gov/upgrades/daq_trig/trigger/reviews/l2_review_020801.html).

### **General Comments**

The Level 2 Decision system for CDF has made excellent progress since the last review on December 7, 2001. Day-to-day monitoring and response to operational issues are handled diligently through a well-organized pager rotation. The switch to a new Magic Bus and TTL arbitration has been implemented successfully. Tremendous progress has been made on Level 2 software, including implementation of PHYSICS trigger table algorithms, distributed development through CVS of algorithm and infrastructure code, version tracking through Tag Sets in the Trigger Database, code optimization, and appointment of designated L2 software coordinators. Important monitoring tools, notably for error detection and detailed timing analysis, provide feedback on the hardware and software performance. The Level 2 decision crate performs well at present machine luminosity (L1A rate  $\sim 10\text{kHz}$  at  $L \sim 2\text{E}31/\text{cm}^2/\text{sec}$ ) and in special rate-test runs (L1A rate  $\sim 20\text{kHz}$ ).

The Level 2 trigger is essential for the success of the CDF experiment; the committee appreciates the amount of work involved in getting the system to this point, and the dedication of the individuals in the Level 2 group. The system has evolved into a reliable and adequate trigger for current needs, but a substantial amount of work remains before CDF can run at full Run 2A luminosity ( $\sim 1\text{E}32$ ) at the highest possible Level 1 rates.

### **Run 2A Performance**

We evaluate the current performance and the expected Run 2A performance against the specifications given in the Run 2 TDR. The specifications are: Level 1 accept rate of  $40\text{ kHz}$ , Level 2 loading takes place in  $10\text{ usec}$ , Level 2 processing takes place in  $10\text{ usec}$ . Level 2 is pipelined so that one event can be processed while the next event

is being loaded. The Level 2 accept rate is 300 Hz, and the specified Level 3 rate is 30 – 50 Hz. The total deadtime is less than 10%.

The Level 2 trigger does not currently meet these specifications, although the system has been able to keep up with the delivered luminosity. Significant improvements have been made since the last review to speed up the Level 2 decision time, including removing or substantially reducing artificial software delays after changing the Magic Bus arbitration from PECL to TTL logic, optimizing the Alpha code, and unpacking data in Alpha only when needed for a trigger algorithm. Currently, loading and processing each take ~25 usec. There are several more improvements that are planned to further reduce the Level 2 decision time.

Although improvements in the trigger system have thus far kept up with Tevatron luminosity increases, expected increases in luminosity in the next few months will lead to increased dead time in the current system. There are three main problem areas: readout/digitization of the SVX, processing time in the SVT, and Level 2 decision time. Each of these should be addressed.

Given that the SVX and SVT times are now the most important factor limiting the Level 1 rate, more attention needs to be given to this matter. Although this point was not explicitly part of the charge of this review committee, we strongly recommend that CDF address these problems and recommend a solution. The Silicon, Trigger, and DAQ (TSI) systems should create a coordinated plan to increase the bandwidth for Level 1 accepts.

Pending issues within the current Level 2 Decision system include the muon interface board, software and firmware modifications, using a single Alpha, low-level data corruption and synchronization errors, and testing/debugging facilities.

Currently, muon triggers account for ~40% of the L2 cross-section. To avoid raising thresholds or prescales, L2 muon information will be necessary to control the L2A rate once the luminosity reaches ~5E31. The lack of a muon interface board in the L2 Decision crate has prevented including muons in the Level 2 trigger so far. We applaud the substantial progress on the L2 muon board that was evident from Fred Neill's talk, and are gratified to see the level of enthusiasm displayed by Fred for commissioning this board. Fred presented a schedule for commissioning which would give us a board working at B0 (at a basic level) by the end of September. Although this schedule seems optimistic, the commissioning effort will benefit from the experience of the Michigan group and the rest of the Level 2 group. We would like to see this effort continue at full speed. We recommend that the commissioning effort for this board be given sufficient resources in terms of manpower and support – at least one additional person should be enlisted for this project. We expect that Fred Neill will be resident at Fermilab during the entire commissioning phase. Weekly progress reports should be given in the Trigger Hardware meeting so that the Trigger group can provide support and expertise as efficiently as possible.

The commissioning of the Muon interface and integration of L2 Muon triggers into the trigger table will require a substantial effort over the coming months. This commissioning will require input from already overburdened experts (e.g. Eric James who is now an Operations Manager and Tom Wright who is responsible for all L2 software). Completion of commissioning on the necessary timescale will require identifying additional people who can work with Fred on aspects such as TrigMon code and L2 algorithms which must proceed in parallel with commissioning the hardware. The Trigger Hardware SPL's should work with Fred to formulate a detailed commissioning plan and to recruit the necessary manpower. This is especially important now that the shutdown is substantially delayed. We may have significant data taking at  $\sim 5E31$  prior to the shutdown in early 2003. At that luminosity, muon information at Level 2 is necessary.

The Level 2 software has made significant advances since the last review. Several people are working together in a well-organized fashion to maintain and improve the Level 2 software. For example, the addition of unpack-on-demand and optimization of the code greatly reduced the Level 2 execution time. It is expected that with further improvements in both software and firmware, the average time to run the L2 algorithms can be less than 10 usec and an average time for the 'decision' (second) stage of Level 2 will be 15 usec. When this goal is achieved, the decision stage will be faster than optimistic estimates for the 'loading' (first) stage, which is currently limited by the  $\sim 27$  usec SVX+SVT latency.

Several improvements to the Alpha firmware are in progress: use of PIO reads while interface boards are sending data (being worked on by D0), improving the L2-TS interface, modifying the FPGA which controls the DMA engine, and improving the VME readout. These improvements will each decrease the Level 2 decision time, and should be pursued at high priority.

The system as designed specifies use of four Alphas to make the trigger decision, although currently only one is being used. It is not clear how much improvement in processing time will be gained by using multiple Alphas, and it is clear that substantial work remains to commission a multiple-Alpha system. In addition, CDF does not have sufficient boards to implement a multi-Alpha system, as four Alphas plus spares would be required, and CDF has only three working in B0. Since the main person who would work on this (Stephen Miller) is already over-committed and he is the only person who can make other improvements to the system in areas where the gain is more evident, we recommend that work on multiple alphas be placed at a lower priority than the proposed firmware modifications.

The situation regarding working spares has improved since the last review, with all interface boards having adequate spares. All boards (except for the Muon Interface board) have the TTL Mbus arbitration mod consisting of blue wires and a glued-on PAL, which should not pose a serious problem as long as sufficient spares and experts are maintained.

The reliability of the system in the past several months is illustrated by the number and reasons for shift-crew calls to the Level 2 general pager. Most of the calls to the Level 2 pager were due to software or infrastructure, with only a handful due to hardware problems in the Level 2 crate. The Level 2 system as a whole has been running reliably and response to problems has been swift and effective, and the pager-carriers are commended for their expertise in dealing with a variety of problems. However, long-term maintenance of the system is a concern, and it was felt that MOU's should formalize the institutional responsibility.

Although the system as a whole functions very well, there is still an ongoing, but low, rate of data errors and errors leading to Decision Timeouts. Decision Timeout errors recoverable by H/R/R occur every 2-3 hours, while more serious errors requiring restarting the run occur less than once per day on average. The rate of Level 2 Decision Timeouts which are caused by problems in the Level 2 Decision crate should be separated from those which are caused by other systems but lead to a Level 2 Decision Timeout or TrigMon errors. The errors due to the Level 2 Decision crate should be categorized as to whether they originate in the Alpha software or firmware, backplane, or interface boards.

Some errors originating in the Alpha may be fixed by planned firmware modifications, although the time required to implement these fixes can be long due to the lack of on-site Alpha expertise and necessity of using CDF as a pulser for the system. Examples of other types of problems are bit errors occurring at a low level, which originate upstream of the Level 2 Decision crate. Debugging facilities must be improved so that problems can be more quickly diagnosed. Ongoing work in this direction is the development of the Pulsar teststand; this effort should be supported. We endorse the ongoing plan to implement an 'error monitor' trigger path, so that many of the events now causing decision timeouts can instead be recorded for offline study without requiring a H/R/R. In addition, we suggest that the messages printed in the alpha's minicom window upon an error should be sent to the Run Control error monitor for later analysis by experts.

The committee would like to comment on the fact that there is a great deal of knowledge concentrated in a few key individuals in the Level 2 group. It is clear that they have done a tremendous amount to build and commission the Level 2 trigger system, and it is necessary for their knowledge to be distributed so that the system can be maintained in the long term. We like the pager rotation as a method for training new experts, and it would be desirable to have more education given to them on all aspects of the system to help ensure the needed transfer of expertise.

To summarize, the Level 2 Decision system is currently performing well, and tremendous progress has been made in the past eight months to improve the Level 2 decision time and to keep the system running stably. We recommend the following:

- a. The most important factor in reducing system deadtime is external to the L2 Decision system; namely the time taken by the SVX and SVT. CDF should

address this by reviewing the factors limiting the L1 accept rate and come up with a plan to alleviate the bottlenecks.

- b. A crucial missing component in the Level 2 decision crate is the muon interface board. Commissioning this board and integrating it into the system should be the highest priority for the Level 2 group. Manpower should be added to this effort.
- c. Alpha firmware and software modifications presented at the review address the second largest cause of deadtime and should continue at a high priority.
- d. Work to run with multiple Alphas should be placed at a significantly lower priority than items a) to c) above. Before any additional work on commissioning multiple Alphas is undertaken, we recommend a study of existing data on contributions to the execution time and particularly the high tail to determine if this will be addressed by multiple processor operation.
- e. Long-term maintenance of the system should be addressed and formalized with MOU's.
- f. Continue to support the development of the Pulsar teststand.
- g. Continue pager rotation scheme; introduce more expert information on all aspects of the system.
- h. It is critical that the first three high-priority items be addressed by separate, non-overlapping teams in order to ensure sufficient manpower.

The Level 2 Decision crate could plausibly then be able to meet CDF's needs for Run 2A.

## **Level 2 for Run 2B**

On the timescale of Run 2B, it is clear that the Alpha boards will be very difficult to maintain. Few spares exist and expertise in repairing the boards will be limited. The DEC Alpha processors and the supporting chip-sets are long obsolete.

Although it is not entirely clear that it is necessary to replace the Alphas solely on the basis of Run 2B bandwidth needs, in order to ensure long-term maintenance of the system we recommend replacing the Alpha boards. Two options were presented at the review: replacing the Alphas with the Beta boards being developed by D0, and a complete replacement of the contents of the Level 2 Decision crate using a Pulsar-based system presented by Ted Liu.

An advantage of the Beta solution is that much of the work has already been done by D0, so it is likely that development time would be shortened relative to starting a similar adapter board project tuned explicitly to CDF's requirements. The Beta boards are designed to be compatible with the existing system architecture, so upgrading could be adiabatic and commissioning the system should be minimally intrusive to the running experiment if the PULSAR test system is operating well and is in place. D0 plans to order their production boards in September, so it would be possible for CDF to order a few boards and make the minor hardware changes required for CDF by hand, in order to evaluate this solution. A potential problem

with the Beta solution is that there are fundamental differences between the needs of CDF and D0, and it is not clear how much work is actually required to adapt the D0 Beta boards to CDF's needs.

The Pulsar board is being prototyped now for use as a teststand for the Level 2 system. The board is designed to accept as input all of the subsystems that feed into the Level 2 Decision crate. In this way, one type of board could replace the six types of custom interface boards now being used. The board uses three large FPGA's where pre-processing of the input data could take place. The Pulsar boards would communicate with a CPU which would be a commercial processor via S-Link, a high-bandwidth optical-to-PCI technology developed at CERN. An advantage of this solution is that it is in principle easier to maintain since there are fewer custom boards in the system. The inherent flexibility and debugging capabilities give confidence that the system could be commissioned quickly in the face of unexpected challenges in the Run 2B environment. A disadvantage of this solution is that it requires the technical risk of replacing the entire Level 2 system. However, this risk is mitigated by the fact that the replacement can happen in parallel with the current working system. Since only a conceptual scheme for using the Pulsar as an upgrade was presented, there are technical details such as data transfer speed and expected deadtime, that are not yet available.

The committee agrees that the Alpha processors should be replaced for Run 2B. There was much discussion regarding the merits and potential problems of the two options presented. The committee concludes that not enough information is available to make a responsible recommendation to CDF as to which upgrade path should be pursued. Therefore we recommend specific steps be taken, and reported on in a future review (timescale = six months):

- a. Further evaluation of the current system is an integral part of making progress toward full Run 2A functionality. Take physics runs with beam at high rate (may require a special trigger table), with a fully loaded backplane (i.e. add a muon board for a single Alpha test and then add one to three additional, but possibly inactive, Alphas) and with all delays removed from the Magic Bus arbitration. Report error rate or other failures. Simulations should be refined and benchmarked (perhaps with test runs) so that deadtime and bottlenecks at Run 2B conditions can be accurately predicted. Including higher occupancy effects and multiple alphas would add to the usefulness of the simulation.
- b. Define the specifications for the Run 2B Level 2 system -- more accurate expectations of occupancies and rates for Run 2B are needed.
- c. Evaluate D0 Beta boards. Purchase two boards (with D0 order in September) and modify with blue wires for CDF-specific signals. Michigan expressed interest in this, and clearly this group has the needed expertise. However, due to the urgent work needed on the current system (which only the Michigan group can do), another group/person should be identified to work on this project in consultation with Michigan.

- d. Evaluate and develop the Pulsar system as a Level 2 replacement system. The commissioning of the Pulsar will begin in the next few months, and provide information on the performance of the board. Measurements should be taken to evaluate the rates, data flow, etc. of the proposed system and compare to the existing system and the Run 2B requirements. Demonstration firmware and software tasks should also be crafted to allow reasonable estimation of the full development schedule. Additional collaboration resources should be immediately identified to help with this project.
- e. Any other proposals for upgrade L2 systems should be presented within this six month time frame.

For all upgrade proposals, a written proposal, including a schedule, cost, and manpower estimate, should be made available. Development time for software and firmware should be addressed.

#### Appendix A: Charge to the Committee

The Level 2 trigger system for CDF has made excellent progress since the last review on December 7, 2001. The new magicbus and TTL arbitration have eliminated the need for the firmware enforced arbitration delays, and will probably serve CDF's needs for the remainder of run 2a with only minor additional hardware modifications. In reading the committee report from that review, it is interesting to see that many of the things discussed there, such as TTL arbitration, and optical splitters for distribution to upper and lower crates, are now implemented.

It is clear that the trigger preparations for Run 2b will depend on the performance of the present system. We request that the committee examine the present L2 system and evaluate it against what CDF will need for Run 2b. This means evaluating the system performance in its present state and understanding what is left to be done to bring the system to its Run 2a baseline level. The evaluation of the performance should include L1 rate capacity, deadtime, reliability, and spares status. The L2 muon board is an important component of the baseline system which is not yet installed in B0. The committee should evaluate the status and the schedule for delivery and commissioning of this board. The committee should determine the schedule for bringing the present system to its "final" configuration. At that point, the committee should determine what, if any, changes will need to be made to the current L2 trigger system to meet the physics bandwidth requirements for run 2b, as specified in the run 2b TDR.

One can imagine that some of the options are:

- 1) Retain the current system, if performance is satisfactory.
- 2) Upgrade the Alpha boards to the D0 L2 Beta.
- 3) Switch to a new L2 decision crate system by replacing the interface boards, magic bus and alphas with a common interface board design and new processors.

Understanding the implications for cost, schedule, and manpower of all of the options above is needed to understand the feasibility of each, so should be included in the report.

#### Appendix B: Committee Membership

Bill Ashmanskas  
Kevin Burkett  
Nathan Eddy  
Bill Foster  
Jane Nachtman (chair)  
Mel Shochet  
Rick Van Berg

#### Appendix C: Level 2 Review Agenda

Introduction	Peter Wilson
System Review	Myron Campbell
Run 2A Baseline / CDF Changes	Greg Feild
Current Status of the L2 Decision Crate	Matt Worcester
Muon Board Status and Schedule	Fred Neill
Software Plan	Tom Wright
Alpha Plan	Stephen Miller
Ultimate Performance Capabilities	Stephen Miller
Physics Requirements for Run 2B on Level 2	
D0 Beta Status	Bob Hirosky
Initial Thoughts of using Pulsar as an Upgrade	Ted Liu